As Cool as a Cucumber: Towards a Corpus of Contemporary Similes in Serbian

Nikola Milošević, Goran Nenadić
School of Computer Science, University of Manchester
Kilburn building, M13 9PL, Manchester, UK,
nikola.milosevic@manchester.ac.uk, g.nenadic@manchester.ac.uk

Abstract

Similes are natural language expressions used to compare unlikely things, where the comparison is not taken literally. They are often used in everyday communication and are an important part of cultural heritage. Having an up-to-date corpus of similes is challenging, as they are constantly coined and/or adapted to the contemporary times. In this paper we present a methodology for semi-automated collection of similes from the world wide web using text mining techniques. We expanded an existing corpus of traditional similes (containing 333 similes) by collecting 446 additional expressions. We, also, explore how crowdsourcing can be used to extract and curate new similes.

Keywords: Phrase modelling, simile extraction, language resource building, crowdsourcing

arXiv:1605.06319v1 [cs.CL] 20 May 2016

1. Introduction

Similes are figures of speech which compares two objects through the use of specific connection words (Harris and others, 2002), but where comparison is not intended to be taken literally (e.g "hladan k'o krastavac" – *as cool as a cucumber*). They are often used as metaphors, but connection words (such as "like" and "as") are used explicitly (Cooper, 1986; Niculae, 2013). In many languages, similes are considered as an important building block of spoken language and are part of cultural heritage. In Serbian, for example, Vuk Karadžić attempted to collect similes, together with other short figures of speech in the XIX century (Karadžić, 1849). While some of the collected similes are still used (e.g. "bela kao sneg" – *white as snow*), some have evolved (e.g. "gladan kao kurjak" – *hungry as a wolf* is now more often appearing as "gladan kao vuk" – *hungry as a wolf*) and some are rarely used and may not even be understood with the current speakers (e.g. "crveni se kao oderano goveče" – *red as skinned ox*). Of course, a number of new similes have emerged. During the XX century, there were attempts to collect proverbs and other short forms from particular geographical areas (Marković, 1979; Jevtić, 1969; Vukanović, 1983; Cvetanović, 1980), but little has been done in updating similes.

To the best of our knowledge, there is no large-scale text mining approach to collecting similes in Serbian, but there have been attempts to collect them in other languages. For example, Veale and Hao (2007) created an approach for automated extraction of explicit similes in English. They modelled them as "X is as P as Y", where X and Y are nouns and P is an adjective. Bin et al. (2012) used similar approach, querying Baidu search engine for Chinese similes. Aasheim (2012) differentiates two types of similes: nominal (*like a/an + noun*) and adjectival (*as + adjective + as + noun*). In Slavic languages, simile expressions are based on the preposition "like" or "as" (*kak* in Russian, *kao* in Serbian, *ako* in Slovak, etc.) and introduce a noun phrase (NP) agreeing in case with the standard of comparison (Rappaport, 1998).

In this paper we explore how we can update the corpus of similes in Serbian using the World Wide Web and crowd-sourcing. The World Wide Web in particular provides people with novel ways to express themselves on a variety of media such as Wikis, question-answer databases, blogs, forums, review-sites and social media. As such, it is an increasingly important resource for exploring linguistic changes and trends, especially as its form is in-between written and spoken languages. Our semi-automated methodology utilises natural language processing, machine learning techniques and manual curation to ensure the validity of extracted data.

2. Method

2.1. Method Overview

The methodology consists of four steps (see Figure 1): firstly we collect the documents for processing; in the second step we process these documents and extract all comparisons; in the third step, we use machine learning to distinguish similes from the other expressions that function as comparisons; in the final step, a human curator reviews the data and corrects any mistakes. Also, users of our crowdsourcing website can propose new similes and help to collect the similes that were not collected by the text mining pipeline; the curator checks these manually.

2.2. Crawler

In order to collect texts from which to mine similes, we have developed a number of crawlers crafted for particular websites. The aim was to download all potentially meaningful text from these websites, but skip contentless parts or parts where similes are unlikely, such as menus, headings and footers. In order to do that, each crawler extracts only text that is inside the div HTML tag with a particular id. Crawlers follow links on each page they visit, but not outside the domain of the website.

We have developed separate crawlers for the following four domains:

- *laguna.rs* - one of the biggest book publishers with abstracts, reviews and parts of the books on their website;

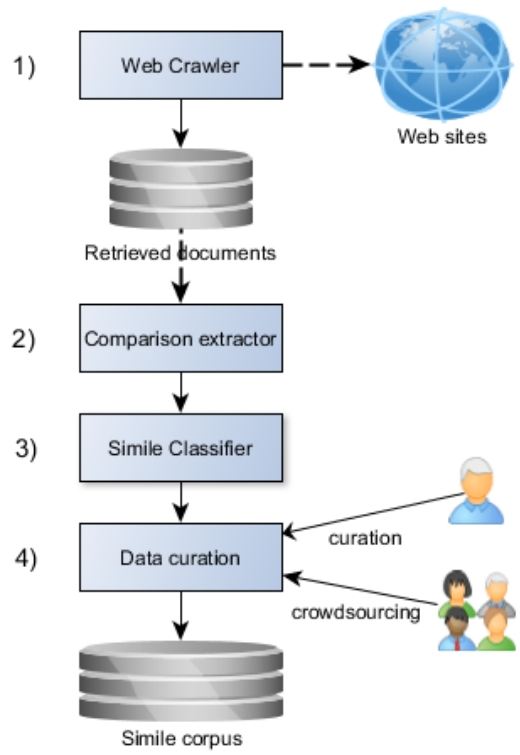

Figure 1: Overview of the proposed methodology: 1) Document retrieval using a web crawler, 2) Comparison extraction, based on part-of-speech tagging, 3) Classification of similes; 4) Data curation with a human curator. Users can also propose missing similes.

- *rastko.rs* - a project that aims to make a digital library of books and articles in Serbian that are considered cultural heritage;
- *burek.com* - a large general public forum on Serbian;
- *tarzanija.com* - a popular blog portal with sarcastic comments on various issues.

For crawling and processing we used Scrapy, an open source Python framework for extracting data from website[1].

2.3. Modelling and extracting candidate similes

The connection word used in similes in Serbian is kao - *like, as*, but it also often appears in shorter forms "k'o" or "ko" (misspelled). There are two main categories of simile in Serbian:

- **Adjectival:** Adjective + Connection word ("kao") + Noun Phrase (e.g. "lep kao cvet" – *beautiful as a flower*)
- **Verbal:** Verb + Connection word ("kao") + Noun Phrase (e.g. "radi kao konj" – *"works like a horse"*)

[1] http://scrapy.org/

Both of these categorise can be modelled by this expression:

$$(V|A|V\ se)\ (kao|ko|k'o)\ (NP) \quad (1)$$

A typical noun phrase consists of a noun together with zero or more dependents of various types, such as determiners, adjectives, adjective phrases, noun adjuncts, prepositional phrases, participle phrases or pronouns (Crystal, 2011). However, currently publicly available part-of-speech taggers and parsers for Serbian can tag only individual words, but not phrases. Because of this, it was not possible to rely on a tagger to identify noun phrases. We therefore modelled *candidate* similes as a verb or adjective followed by explicit use of connection words "kao", "k'o" or "ko", followed by one or more adjectives and terminated with a noun:

$$(V|A|V\ se)\ (kao|ko|k'o)\ ((A|N)*)\ (N) \quad (2)$$

The noun phrase model that we used models the majority of commonly used noun phrases in Serbian. However, noun phrases and similes often contain prepositions. e.g. "smoren kao zmaj u vatrogasnoj stanici" (*bored like a dragon at a fire station*). Our model will currently pick up "smoren kao zmaj" (*bored like a dragon*), which is also a simile used quite frequently, but we would rely on manual curation to extend it with the prepositional phrase.

In the next step of the pipeline, we use retrieved textual documents to extract candidate comparisons. We used a part of speech (POS) model for Croatian and Serbian (Agić et al., 2013) and plugged it into the HunPos tagger (Halácsy et al., 2007). We used the model 2 for matching similes.

2.4. Simile classifier

After a review of initial outputs, we noted a large number of false positives. The model would find similes such as "radi kao konj" (*works like a horse*), but it would also extract expressions like "radi kao pravnik" (*works as a lawyer*), which is not a comparison, but rather specifies someone's profession. There is no lexical feature that would give a cue that these two phrases are in any way different and only semantics (and context) would differentiate the similes from other expressions. However, current publicly available tools for Serbian do not process phrases at the semantic level and thus we decided to rely on a classifier and a final curator's manual review of the phrases.

From the initial data, we have created a sufficient number of true positive and false positive examples that could be used for machine learning. This data was used to create a machine-learning based classifier that can distinguish true similes from the other phrases that have same lexical characteristics. The dataset contained 300 examples of true positive and 300 examples of false positive similes to train and test our algorithm. For the features, we used

- the whole simile phrase (for illustration we will use "radi kao konj" – *"works like a horse"*),
- the stemmed phrase ("rad ka konj"),
- left side of the phrase that is before the connection word ("radi" – *works*),

- left side stemmed ("rad"),
- the part of phrase that is on the right from the connection word ("konj" – *horse*), and
- the stemmed right part ("konj").

For stemming we used stemmer for Serbian (Milošević, 2012), that was ported for Python[2]. We experimented with Multinominal Naive Bayes, Random Forests and Support Vector Machines (SVM) with a polynomial kernel that uses sequential minimal optimisation. For machine learning we used Weka toolkit (Hall et al., 2009). The models are downloadable from the corpus website[3].

2.5. Curation interface

We have created a web portal[4] for viewing, searching and editing the dataset collected from the previous steps. The portal has the following curation features:

- **Viewing of all currently curated similes**, sorted alphabetically
- **Searching for simile.** Searches are performed by using the stemmer for Serbian, so the similes that are not in exactly the same form in the database can be retrieved. Since Serbian is highly inflectional, the aim was to store only one canonical instance (single inflection) of each simile. For example "beo kao sneg" (m), "bela kao sneg" (f) and "belo kao sneg" (n) (*"white as a snow"* in three different grammatical genders - masculine, feminine, neuter), will be treated by our search algorithm as the expression. Search interface can be seen in Figure 2.
- **Adding a new simile to the database.** Users of the website are able to manually add similes that are rare or missing from the corpus. If the user wants to add simile that already exists in the database, he/she will be notified. Added simile won't be visible on the website until curator approves it.
- **Curation interface.** Curators can login to the website and perform curation task, such as approving, rejecting or editing similes.

3. Results

The crawlers downloaded 40,239 documents from the web (See Table 1). Using semi-automated workflow, we extracted potential similes and manually reviewed them. A total of 446 true similes were finally harvested. A manual review of a data also revealed that there were around 5,000 false positives. Our classification approach proved to be useful in filtering some of these out - the results are presented in Table 2.

[2] https://github.com/nikolamilosevic86/SerbianStemmer

[3] http://ezbirka.starisloveni.com/download.html

[4] http://ezbirka.starisloveni.com

Website	Number of documents
burek.com	22,925
laguna.rs	9,574
tarzanija.com	4,792
rastko.rs	2,947
Total	40,239

Table 1: Number of documents in dataset obtained by crawlers

Algorithm	Precision	Recall	F-Measure
Random Forests	0.783	0.732	0.756
Naive Bayes	0.782	0.773	0.777
SVM	0.804	0.791	0.797

Table 2: Results of machine learning based classification of similes

We merged the similes obtained from the web with the simile corpus published by Vuk Karadžić (Karadžić, 1849), who collected 333 similes. It is interesting that there was only a small overlap between the two simile datasets. Overlap between these two corpora can be seen in Table 4 The similes collected by Karadžić were generally more complex, which could be due to the termination rule on noun in our model, so the model might be missing some complex similes. Although, this is an obvious limitation of the system, collection of complex similes will be part of the future work, and also will rely on crowdsourcing interface. At the moment of writing this paper, we had 852 approved similes in our corpus.

Corpus	Number of similes
WWW Corpus	446
Karadžić's Corpus	333
Intersection between Karadžić's and WWW Corpus	9
Manually added similes	80
Total similes	852

Table 3: Number of similes obtained from World Wide Web, Karadžić's corpus and initial crowdsourcing. Also, number of similes that appear in both WWW and Karadžić's corpus is presented

Original simile	Translation
beo kao sneg	white as a snow
brz kao zec	fast like a rabbit
crven kao vampir	red like vampire
crven kao krv	red like blood
miran kao ovca	calm as a sheep
lako kao pero	light as a feather
ljut kao ris	angry like a lynx
slatko kao med	sweet like honey
tvrd kao kamen	heavy like a stone

Table 4: Intersection between similes found on the WWW and in Karadžić's collection

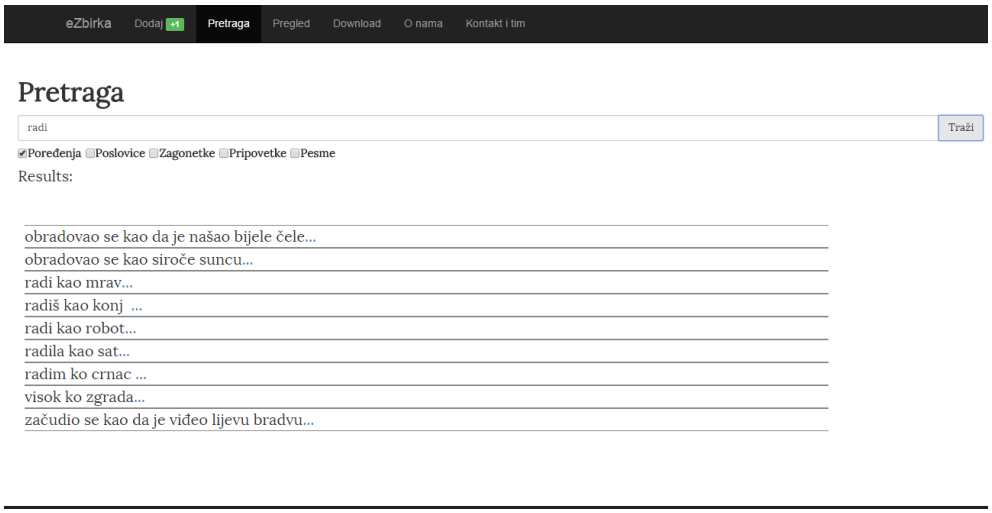

Figure 2: Simile search interface

4. Summary

In this paper we present initial work towards building a corpus of similes in Serbian and a methodology for a semi-automated creation of simile corpus. We have kick-started the process by processing a set of web sites that are likely to have a number of similes. The current corpus, with 852 similes is, to the best of our knowledge, the largest simile corpus in Serbian. We believe that, by mining more web pages for similes and by the use of crowdsourcing, the corpus will grow and remain up to date. Still, the current methodology will need to be improved through the use of semantic resources (e.g. WordNet) and contextualisation.

The International Standard Language Resource Number (ISLRN) of the Serbian Simile Corpus is *534-837-568-404-3*. It is freely available at `http://www.ezbirka.starisloveni.com/`. The code of the simile extractor can be found on GitHub[5].

5. References

Aasheim, I. (2012). *A contrastive study of similes in English and Norwegian*. UNIVERSITY OF OSLO,A master thesis presented to the Department of Literature, Area Studies and European Languages.

Agić, Ž., Ljubešić, N., and Merkler, D. (2013). Lemmatization and morphosyntactic tagging of croatian and serbian. In *Proceedings of ACL*.

Cooper, D. E. (1986). *Metaphor (Aristotelian Society Series)*. Blackwell Pub.

Crystal, D. (2011). *Dictionary of linguistics and phonetics*, volume 30. John Wiley & Sons.

Cvetanović, V. (1980). *Narodna književnost Srba na Kosovu. 7, Narodne poslovice*. Jedinstvo, Priština.

Halácsy, P., Kornai, A., and Oravecz, C. (2007). Hunpos: an open source trigram tagger. In *Proceedings of the 45th annual meeting of the ACL on interactive poster and demonstration sessions*, pages 209–212. Association for Computational Linguistics.

Hall, M., Frank, E., Holmes, G., Pfahringer, B., Reutemann, P., and Witten, I. H. (2009). The weka data mining software: an update. *ACM SIGKDD explorations newsletter*, 11(1):10–18.

Harris, R. et al. (2002). *A handbook of rhetorical devices*. Virtual Salt.

Jevtić, M. (1969). *Narodne poslovice i uzrecice iz Podrinja*. Godišnjak istorijskog muzeja u Šabcu.

Karadžić, V. S. (1849). *Srpske narodne poslovice i druge različne kao i one u običaj uzete riječi*.

Li, B., Kuang, H., Zhang, Y., Chen, J., and Tang, X. (2012). Using similes to extract basic sentiments across languages. In *Web Information Systems and Mining - International Conference, (WISM) 2012, Chengdu, China, October 26-28, 2012. Proceedings*, pages 536–542.

Marković, R. (1979). *Metaforisane i nove poslovice Šumadije*. Narodni muzej, Čacak.

Milošević, N. (2012). Stemmer for serbian language. *arXiv preprint arXiv:1209.4471*.

Niculae, V. (2013). Comparison pattern matching and creative simile recognition. In *Proceedings of the Joint Symposium on Semantic Processing*.

Rappaport, G. (1998). The slavic noun phrase. *Position paper for Comparative Slavic Morphosyntax]. Available: http://www. indiana. edu/˜ slavconf/linguistics/download. html [1999, Jan. 4]*.

Veale, T. and Hao, Y. (2007). Making lexical ontologies functional and context-sensitive. In *ANNUAL MEETING-ASSOCIATION FOR COMPUTATIONAL LINGUISTICS*, volume 45, page 57.

Vukanović, T. (1983). *Vranjanske narodne poslovice*. Nova Jugoslavija, Vranje.

[5] `https://github.com/nikolamilosevic86/SerbianComparisonExtractor`